



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 539 106 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention
of the grant of the patent:
21.08.2002 Bulletin 2002/34

(51) Int Cl.7: G06F 17/30, G06K 9/20

(21) Application number: 92309445.2

(22) Date of filing: 16.10.1992

(54) Electronic information delivery system

Elektronisches Datenliefersystem

Système électronique de livraison d'information

(84) Designated Contracting States:
FR GB

(30) Priority: 24.10.1991 US 782842

(43) Date of publication of application:
28.04.1993 Bulletin 1993/17

(73) Proprietor: AT&T Corp.
New York, NY 10013-2412 (US)

(72) Inventors:

- Fox, David S.
Summit, New Jersey 07901 (US)
- Jagadish, Hosagrahar V.
Berkeley Heights, New Jersey 07922 (JP)
- O'Gorman, Lawrence
Madison, New Jersey 07940 (US)
- Story, Guy A.
NY, New York 10012 (US)

(74) Representative:
Watts, Christopher Malcolm Kelway, Dr. et al
Lucent Technologies (UK) Ltd,
5 Mornington Road
Woodford Green Essex, IG8 0TU (GB)

(56) References cited:

- TRANSACTIONS OF THE INSTITUTE OF ELECTRONICS AND COMMUNICATION ENGINEERS OF JAPAN, SECTION E, vol.E72, no.6, June 1989, TOKYO JP pages 771 - 781 YAMAMOTO A. ET AL. 'Extraction of Object Features and Its Application to Image Retrieval'
- SYSTEMS & COMPUTERS IN JAPAN, vol.21, no.11, 1990, NEW YORK US pages 33 - 46 KATO T. ET AL. 'TRADEMARK: Multimedia Image Database System with Intelligent Human Interface'
- PATENT ABSTRACTS OF JAPAN vol. 15, no. 131 (P-1186) 29 March 1991 & JP-A-03 012 889 (MATSUSHITA ELECTRIC IND. CO. LTD.) 21 January 1991
- PROC. 1986 IEEE FALL JOINT COMPUTER CONF., 2 November 1986, DALLAS, USA pages 87 - 96 SRIHARI S.N. 'Document Image Understanding'
- IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, vol.10, no.6, November 1988, NEW YORK, US pages 910 - 918 FLETCHER L.A. ET AL. 'A Robust Algorithm for Text String Separation from Mixed Text/Graphics Images'

EP 0 539 106 B1

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

[0001] This invention relates to electronic libraries and more specifically to electronic storage and retrieval of information. More generally, this invention relates to storing documents and other visual information segments in a computer, and retrieving those documents in response to user queries.

[0002] For decades the electronic library has been the holy grail for far-thinking librarians, writers, and computer scientists. Examples include the "world information monopoly", presented in 1936 by H.G. Wells in *World Brain*; MEX-MEX, described in 1945 by Vannevar Bush in his classic article *As We May Think*; and others.

[0003] Full-text databases such as NEXIS and NewsNet are now available, and since they do provide text information they may be considered "libraries". With NEXIS, for example, the wealth of information that can be gathered in a very short time is staggering, and this makes these databases very powerful and valuable resources. However, they lack the ability to store, search, retrieve and display pictures, graphs, etc., and they do not offer a comfortable environment for browsing.

[0004] It is important to realize that ASCII representations of text, such as found in NEXIS do not convey all of the information that an original printed text does. Disregarding possible transcription errors, some information simply cannot be conveyed when the original printed text is discarded. In addition to the aforementioned pictures, accenting and highlighting are lost, mathematical equations are almost impossible to comprehend, unusual symbols cannot be represented, etc.

[0005] Also, a fair amount of information is contained in the position of text on the page (such as in a business letter), and that information is lost in systems such as NEXIS. The written word has had been around for centuries, and writing formats have been adopted over the years that are easily recognized and understood by readers. These formats have been honed to convey information efficiently. Thus, a mere space can in some circumstances communicate as much information as an undecipherable scribble, or a whole sentence. For example, a substantial space in the beginning of a line and before a sentence is recognized as a paragraph delimiter and tells the reader that a new thought is about to be addressed. A scribbling at the end of a letter indicates that the letter was signed by the sender, even if not all of the characters in the scribbling are recognizable.

[0006] Additionally, those who regularly read particular types of documents develop a facility to jump to the most important portion of the document simply based on position of the information or other indicia (such as the largest paragraph of a memo, shortest paragraph of the memo, the sentence that is underlined, etc.) That, too, is mostly lost in systems such as NEXIS.

[0007] Lastly, the fact that people are simply comfortable with the familiar formats of newspapers, magazines and books should not be underestimated and that familiarity is mostly lost in data base systems such as NEXIS.

[0008] Commercial image databases on CD-ROM, such as those recently introduced by UMI, are closer to an electronic library, in that they provide images of the stored pages. This permits the stored images to contain text and pictures. However, these systems are very limited in their search and retrieval capabilities because they require a manual abstracting and indexing of the stored images to provide a key word search capability.

[0009] In the optical character recognition (OCR) art, it is now possible to purchase an OCR system that can scan a page of text and identify the printed ASCII characters contained therein, as well as identify the font and size of those characters. Typically, the OCR systems are processor controlled and the (more advanced) programs that implement the OCR recognition algorithms consult a dictionary when a letter is difficult to recognize. The end result of these OCR systems is that a scanned page of text is converted to ASCII form, as best as the program can, and the ASCII form is stored in the system's memory. Upon request, the ASCII text (as good or as bad as it may be) is displayed to the user. The scanned image is not kept.

[0010] Even with the availability of all of these diverse capabilities, there is still not a single system that approaches the functionality of a conventional library.

[0011] A. Yamamoto et al., in The Transactions of the IEICE, vol E72, no [6], June 1989, pages 771-781 disclose an image retrieval method based on object features which contain 2-D information of original images and correspond to human conceptions in terms of image composition.

[0012] T Kato et al., in Systems and Computers in Japan, vol 21, no [11], 1990, pages 33-46, disclose an experimental multimedia database system being developed, which includes image processing facilities as the primitive data operations.

[0013] S.N. Srihari, in the 1986 Proceedings of the Fall Joint Computer Conference, November 2-6, 1986, Dallas, Texas, USA, pages 87-96, discusses the problem of document understanding, which is a goal-oriented problem involving detecting and interpreting different blocks and coordinating the interpretations to achieve an end result.

[0014] Methods of storing and accessing according to the invention are as set out in the independent claims.

[0015] This invention provides the means for realizing an electronic library that very closely emulates the interaction modes of a physical library. Specifically, the electronic library of this invention maintains an electronically searchable image of all information that it maintains, but, it delivers to the user an audio visual image in response to a user's

request. In accordance with the principles of this invention, the electronic library comprises an electronic user interface, such as a computer screen, a speaker, a mouse and/or a keyboard; a processor for handling communication with the user and for responding to user requests; and a data store. The data that is stored is derived from segments of information which were scanned, processed and stored in the system. It is the scanned segments of information (or portions thereof) which, upon request, are provided to the user. Those segments may be images of journals, newspapers, letters, magazines, maps, graphs, etc., and they may also be digitized segments of speech, music and other audio sources. In addition to the stored segments of information that are displayed to the user upon request, translated versions of the same data are also maintained in the data store. The translated versions contain the immediately translatable version of the displayable information and processed information that forms various aggregations of the displayable information. This processing imposes a syntactically logical structure on the displayable information. It is the translated version of the data that forms the electronically searchable source of data.

FIG. 1 illustrates in very broad strokes the hardware arrangement for realizing an electronic library, and depicts some of the data structures within memory 102;

FIG. 2 presents a flow diagram of a search, retrieve and display process in accordance with the principles of this invention;

FIG. 3 illustrates in greater detail the three planes of information which are associated with each displayable segment;

FIG. 4 outlines the processes for developing the information in plane 2 and in plane 3;

FIG. 5 presents a flowchart of texture analysis; and

FIG. 6 presents a more detailed flow diagram of the processes associated with the plane 3 information.

[0016] FIG. 1 presents an overview of an electronic library system embodying the principles of this invention. Element 100 is the computer screen-keyboard-speaker-printer arrangement that forms the user interface. A user can direct instructions and search queries to the system via the keyboard (or the mouse), and the system responds by either displaying information on the screen or printing it (when the data is visual), or outputting it through the speaker (when the data is aural). Element 100 is connected to processor 101 which interacts with memory 102 and memory 102 contains one or more databases of scanned and digitized segments. Blocks 103 and 104 represent two segments which are stored in memory 102.

[0017] It should be pointed out that the types of information that are stored in memory 102 can be quite diverse. The information may be all text, akin to the information stored in the NEXIS database; it may be text co-mingled with pictures, such as magazine articles; it may be primarily picture information, such as charts, graphs, photographs, etc; and it can even be speech or music. Also, there can be more than one database that is stored in memory 102, and the databases do not have to store similar types of data.

[0018] FIG. 1 depicts only the two digitized segments 103 and 104, and they are shown side by side. While this suggests a plurality of segments in memory 102, it does not describe how the segments are stored within the memory or, indeed, what information they represent.

[0019] The information stored in a particular database might advantageously be stored in a hierarchical structure. For example, one may wish to create a database of technical journals in a particular field of interest. For such a database, at the highest hierarchical level, specific journals are identified. At the next (lower) level, different issues of a selected journal are identified. In the following level, different articles in a selected issue are identified. In a still following level, different pages of a selected article are identified, and in the lowest level, perhaps, different paragraphs are identified.

[0020] The term "segment" in this disclosure assumes a meaning that comports with the context in which it is used. When seeking to select a particular article from a collection of articles, "segment" is an article. When searching for a particular page within an article, "segment" is a page. Most of the time, however, the term "segment" refers to a quantum of information that is stored in memory 102 and which can be (or is intended to be) provided to the user as a block. When the information is an image, that may mean the information fits on the screen of arrangement 100.

[0021] Returning to FIG. 1, in accordance with the principles of this invention three planes of information are associated with each digitized segment. The first plane contains the digitized representation of the scanned segment itself (e.g. blocks 103 and 104), the second plane contains elemental information that is found in the digitized image (this is shown by blocks 113 and 114), and the third plane contains macro information which identifies groupings of elemental information (this is shown by blocks 123 and 124). When the digitized and scanned segment is an image from a magazine, the elemental information entities of the second plane are letters, lines, symbols, and the like. The macro elements in the third plane are logical groupings such as a title sentence, the author's name, a date, a picture block, and the like. The "information" in the second and third planes form the set of translated information. That is, the information in blocks 113 and 123 contain translations, or transformations of the information in block 103.

[0022] FIG. 2 presents a general flow diagram of an information retrieval process that derives information from a

database setup as the one described in connection with FIG. 1. In block 200, the user enters a query into the system. That query can be formulated by selecting an icon on the computer's screen, by actually typing in the query. This step is completely conventional. Having received a search request, block 102 searches in the translated planes, (e.g., through the information in blocks 113, 114, 123 and 124) for an information segment that matches the search request.

5 The type of segment searched for can be specified by the user as part of the search query, or it may be some default segment type that, perhaps, depends on the type of search specified in the query.

[0023] Another way a search query can be formulated is by pointing to, and highlighting, a region on the screen which displays the image. The mechanics of identifying a region are well known. It can be found, for example, in many of the word processing programs that are commercially available, where a cursor may be pointed to a line of text and "dragged" to highlight a portion of a line or a plurality of lines. One difference, however, is that the image shown on the screen is a rendering of stored information, and it is the stored information that is being "highlighted", or linked. In the context of this invention, it is the scanned image that is displayed, it is the scanned image portion that is highlighted and a link to the translated image is identified. The query is executed on what the translated image contains.

[0024] The search performed by block 201 results in any number of "hits". If that number is 1 or greater, block 202 stores a pointer to the identified segments of the first plane and to the translated segments of the second and third planes. Thereafter, a display step is carried out by blocks 204 and 205. Specifically, when block 202 contains more than one "hit", block 204 displays one of the digitized segments pointed-to by block 202 and waits for a user input. Block 205 responds to the user's direction. When a user directs the display of the next search result, control passes to block 204 through line 206, directing block 204 to display a different one of the pointed-to digitized segments. When

20 the user requests a new search, control returns to block 200 via line 207.

[0025] FIG. 3 presents an example of the three planes of video information of a page from, say, a notebook. Plane 1, shown as frame 105, contains the letters "ABC" -- that being the text on the page -- two diagonal lines slightly below and to the right of "ABC", the equation "E = mc²" below the diagonal lines and in the center of the page and, lastly, a grey-scale picture below the equation. It should be appreciated that although frame 105 in FIG. 3 is depicted in a manner that is recognizable to the human eye, in reality that information is stored in memory 102 as a block of individual pixels having specified darkness levels (alternatively, to reduce storage the pixels can be encoded in any one of well known techniques).

[0026] Plane 2, with data that relates to the data of block 105 and which is marked 106, illustrates one manner by which the information contained in the digitized image of 105 may be stored in the form of elemental information entities that are contained in the image. It may be noted that one of the elemental information entities is an "unrecognized box". This entity can encompass not only pictures but other markings on the page, such as unrecognized letters, symbols, scribbles, doodles, etc. In some embodiments, this elemental information entity may also encompass all handwritten letters (such as signatures).

[0027] Plane 3, with data that relates to planes 1 and 2 and which is marked 107, contains the macro elements, or blocks, that are found in the image. In FIG. 3 frame 107 contains only four entries: one for the text, one for the diagonal lines, one for the equation, and one for the picture. Table 108 is the table of pointers that ties the logical page blocks of frame 107 to the elemental information entities of frame 106 and to the digitized segment of frame 105.

[0028] The actual programs for implementing the search scheme described above in connection with FIGS. 2 and 3 are completely conventional. Almost any commercial database manager program can be augmented (by adding appropriate program modules) to incorporate that aspect of this invention which identifies a translated segment but displays the digitized segment that is associated with the translated segment.

[0029] The more challenging task is to create the translated planes from the raw scanned data. This process is outlined in FIG. 4 which comprises two parallel paths which diverge from the scanned image output of block 301: one path contains blocks 302 and 303, and the other path contains block 304.

[0030] Image segmentation block 302 identifies areas in the scanned image output of block 301 which cannot be recognized by a conventional optical character recognizer. The purpose of block 302 is to eliminate from consideration by the following OCR block (303) those areas of the image that do not contain information that can be identified by the OCR block. There is a need to find those areas anyway, and there is no sense to burden the OCR block with analyses that will not be fruitful. In the case of FIG. 3, the task of block 302 is to identify the lines and the grey-scale picture in plane 105. This is accomplished with texture analysis of the image to determine types of regions and classifying them as: blank, text, line diagram, equation (or symbolic line of non-ASCII elements), line segmentors, binary picture, dithered picture, grey-scale picture and color picture.

[0031] The texture analysis can be performed on a gray-scale image as illustrated in FIG. 5. In block 400, the image is divided into contiguous, non-overlapping, windows of a size that is slightly larger than the most commonly occurring character size. In block 410 the pixels in each window are examined to determine the entropy (measure of disorder, or variation) in the window. Regions whose entropy measure is low are labeled as binary (suggesting that the window contains either a character or a portion of a line drawing) windows, and whose entropy measure is larger are labeled as gray-scale picture windows. In block 420 the label attached by block 410 to each window is reviewed, based on the

EP 0 539 106 B1

8 nearest neighbors of each window, and corrected if necessary. Individual windows are corrected so neighboring windows that have the same label form regions with an expected shape and size that is appropriate to the particular page being analyzed (usually rectangular), known a priori by domain information.

[0032] The binary labeled windows (and regions) are directed to block 430 where the input is binarized; i.e., set to 1 or zero by thresholding with respect to a value intermediate between high and low gray-scale values.

[0033] The output of block 430 is applied to block 440 where the binary label is refined as follows:

- if there is a high percentage of 1-valued pixels in a window, with the 8 neighbors being 0-valued, then the window is labeled as a dithered window;
- if there is only 1 or a few connected (neighboring) windows of 1-values, with a proportion of 1- to 0-valued pixels being about 1/16 to 1/8, then the window is labeled as a text window;
- if there is only 1 or a few connected windows of 1-values, with a proportion of 1- to 0-valued pixels being less than about 1/16, then the window is labeled as a line graphics window;
- if there are no 1-valued pixels in the window, then window is labeled as an empty window;
- if there are only 1-valued pixels, then the window is labeled as a binary picture window (black);

[0034] The output of block 440 is applied to block 450 where the label attached by block 440 to each window is reviewed, based on the 8 nearest neighbors of each window, and corrected if necessary. As in block 420, individual windows are corrected so neighboring windows that have the same label form regions with an expected shape and size that is appropriate to the particular page being analyzed (usually rectangular), known a priori by domain information.

Lastly, the gray scale output of block 420 and the output of block 440 are combined and formatted in block 460 to form the "plane 2" output of block 302, and in block 470 to form the mask for OCR block 303.

[0035] Thus, block 302 identifies the unrecognized segments in the scanned image and develops a template that instructs OCR block 303 to ignore certain areas of the image.

[0036] Block 303 which follows block 302 is a conventional OCR block (e.g., Calera RS-9000) and it identifies the characters of frame 105. The combined output of image segmentation block 302 and OCR block 303 forms frame 106. The output of block 303 is the plane 2 information.

[0037] FIG. 6 presents a more detailed flow chart of the process carried out in block 304 of FIG. 3. Block 305 accepts the scanned image information of block 301 and analyzes that information to identify connected components. The set of connected components developed by block 305 are applied to block 306 which determines the K-nearest neighbors of each block. Results of the analysis performed in block 306 and provided to block 307 merges image components to create logical blocks. Lastly, the logical blocks developed by block 307 are analyzed through a parsing process in block 308 to obtain a syntactic segmentation of the information contained in the scanned image. The syntactic segmentation is guided by information provided by block 309, which is derived from *a priori* knowledge of the format of the scanned image.

[0038] The task of determining connected components on the scanned image (i.e., the task of block 305) can be carried out as follows. First, consider every pixel in the image, and for every pixel that has value "ON" (i.e., it is darker than a preselected darkness level), determine whether 1 or more of its 8 closest pixels (N,S,E,W,NW,NE,SW,SE) has value OFF. If so, label the center pixel as a contour pixel. After this, link the contours into chains by first searching the image pixels in any sequential (row-column or column-row) order in the neighborhood of the found contour pixel until another contour pixel is found. Once a contour pixel is found, follow the contour from one contour pixel to a neighboring contour pixel, erase (set to non-contour value) each such pixel and also store the location of each contour pixel in a vector labeled by a distinct connected component number (index). That vector designates a connected component, e.g., a character. Continue populating that vector until there are no neighboring contour pixels left. Thereafter, find another contour pixel and begin identifying a new connected component. Repeat the process of identifying connected components until no contour pixels remain. Lastly, determine the centroid of each connected component, and store the location (x and y position) of the centroid for each of the connected components. The x position of the centroid is determined by adding the horizontal positions of all of the contour pixels in the connected component and dividing the sum by the number of such pixels. Similarly, the y position of the centroid is determined by adding the vertical positions of all of the contour pixels in the connected component and dividing the sum by the number of such pixels.

[0039] The k-nearest neighbor analysis (block 306) can be carried out by choosing a value for K (typically 3, 4 or 5), and for each connected component, finding the minimum Euclidean distance from the centroid of the connected component to the centroids of each of other components. The nearest K neighbors are identified and their indexes are stored in association with the connected component, along with distance and angle to each. The result of this process

EP 0 539 106 B1

is a table that may have the following format:

5	Comp i	Neighbor 1			Neighbor 2			Neighbor 3		
		index	dist. from i (pels)	angle from i (°)	index	dist. from i (pels)	angle from i (°)	index	dist. from i (pels)	angle from i (°)
	3	12	30	5	17	52	91	18	28	4

10 [0040] To merge connected components to create segment blocks (block 307), one needs to first determine the skew angle of the image, the inter-character spacing, the inter-word spacing and the inter-line spacing.

[0041] The skew angle is determined by finding the peak angle of all neighbor pairs from a histogram of these angles.

[0042] The inter-character spacing is determined by grouping the pairs in distance range groups. Pairs that correspond to adjacent characters will have a small distance within some range and they will form the most populous group. Averaging the distances in this group yields the nominal inter-character spacing. The group of pairs that are within the range of a slightly larger average distance are the pairs of characters at the end of one word and the beginning of the next word. Averaging the distances in that group yields the nominal inter-word spacing.

[0043] The inter-line spacing is determined by identifying all pairs where the angle relative to the skew angle is greater (in magnitude) than 45° ; and for all such pairs, finding the most frequent average distance. Finding the most frequent average distance means observing that the pair distances can be grouped into distance ranges, identifying the group that contains the largest number of pairs and computing the average pair distance for that group. This is the inter-line spacing.

[0044] From the above it is easily appreciated that "words" are groups of connected component pairs whose angles are within 45° of the skew angle and whose pair distances are within a chosen tolerance of the inter-character spacing. "Lines" of text are groups of connected component pairs whose angles are within 45° of the skew angle and whose distances are within a chosen tolerance of the inter-word spacing. "Blocks of text" are lines of text whose average inter-line spacings are within a chosen tolerance of the inter-line spacing.

[0045] As depicted in FIG. 3, the "words", "lines", and "blocks" are stored in translated plane 3 with pointers to their connected components, and spatial coordinate information.

[0046] The segment blocks identified in block 307 are parsed in block 308 to create logical blocks of the segment. To do its parsing, block 308 employs "domain-dependent" information that is provided by block 309. The information of block 309 is supplied by the user. This information may state, for example (for correspondence letters), that *"the date is found at about 2 inches from the top of the pages, it is a block of text that is shorter than 2 inches and it has at least 4 inches of white space to its left"*. It is also likely to state that *"the subject is a block of text that is indented, is below a line that begins with 'Dear', and it starts with 'Re:'"*, etc. This is the information that describes the characteristics of a page that make it a "correspondence letter".

[0047] Once the information is parsed and block 107 of FIG. 2 is populated with the information that specifies the logical blocks, a linking must take place between the different elements of the three planes of information. This is accomplished in a straight forward manner as outlined below, and the results placed in table 108.

[0048] The initial information is: coordinates of each pixel on plane 1; coordinates of bounding boxes or contours of the connected components, which are the elemental entities of plane 2 discussed above; and coordinates of blocks on plane 3. Therefore, planes 2 and 3 are linked to plane 1. To link planes 2 and 3, one needs to merely test each elemental entity in plane 2 for overlap between the elemental entity box and a block in plane 3. If the elemental entity overlaps a block, a pointer is created from the elemental entity to the block, and vice-versa. These pointers are incorporated in table 108 shown in FIG. 3.

[0049] As indicated above, from the standpoint of the mechanics of performing the search, searching for particular information within the FIG. 1 arrangement is fairly conventional. That is, various database manager programs can easily be tailored to effect searching in blocks 106 or and/or 107. However, because of the particular structure of the FIG. 1 arrangement, the overall effect is very powerful. For example, aside from the capabilities made available by the hierarchical structure of the databases, the parsing carried out by the process of block 308 creates a wealth of information that may be used *during a search and retrieve session*. In a database of correspondence letters, for example, the domain-dependent information in block 309 is likely to identify the addressee of a letter, the date of the letter, the "Re:" line of the letter, the sender, and perhaps the sender's company. Consequently, searching can be performed on any one of those categories. The identity of categories is, in effect, automated by the parsing process and the search categories come naturally. That is, the categories that are selectable (and searchable) originate from the domain-dependent information of block 309 and, hence, are easily changeable from one data base to another. These categories can be easily displayed on the screen via an icon arrangement, and searching in a category can be specified by merely pointing to an icon. Search of categories can also be performed through the search window (see element 100 in FIG.

1) by specifying the search in a manner akin to the specification of the search in the NEXIS database.

[0050] The following three examples illustrate just some of the power in the search and display capabilities of this invention.

[0051] As a first example, it is not uncommon for a page of text in a technical article to contain references to information that is not found on the very same page. One example is references to materials that are identified by the author as being relevant to the subject at hand ("references"). In many magazines the format for indicating that a reference is being identified is very specific (e.g., initials of the author and a number, such as [JPM88]). This format forms a piece of information that can be included in block 309. When the domain information specifies this format and an instance of the format is found in the scanned page by block 308, in accordance with the principles of this invention a link is provided between the string, e.g., [JPM88], and the image segment which contains the details of the reference identified by [JPM88]. With this link in place, when a user obtains a display of a text which contains [JPM88] and the user highlights this string, the image segment which contains the details of the reference is retrieved from the database and displayed, preferably in a second window, on the computer display.

[0052] Another instance of a reference to information that is not contained in the displayed page is often found with references that relate to figures, plots, tables, pictures, etc. Again, this invention permits the domain-information to form a link which associates a reference such as "FIG. 3" with another image segment that is stored in the computer which represents FIG. 3 (i.e. which has a "title" that comprises the string "FIG. 3").

[0053] The above-described capability introduces a very powerful tool for computerized browsing which is not found in any prior art computerized system but which is easily realized and often used when a user reads a printed version of an article.

[0054] As a second example, over and above having the ability to refer to specific portions of the same article that are referenced by indications such as "FIG. 3" and "[JPM88]", the FIG. 1 arrangement offers the capability to actually call up the indicated reference (i.e., the reference abbreviated by [JPM88]) or any other reference that the user may want to view immediately (or concurrently) either in the same window or in a separate window of the computer screen.

[0055] As a third example, even when the OCR and associated processes do not faithfully convert, or translate, the scanned image, it is still possible to identify entries based on a key word that is slightly corrupted and then permit the user to correct the translation. It is also possible, with the FIG. 1 arrangement, for the user to observe that some input word is either poorly scanned or, perhaps, misspelled in the original, and correct it. The correction is effected by the user highlighting the image portion of his choice, whereupon the translated version of the highlighted portion is shown on a separate screen. That translated portion can then be edited for future reference and use.

[0056] It may be noted that in describing the OCR process (303) no mention was made of the specific OCR process that is employed, other than suggesting that a conventional one may be used. In fact, we employ an OCR process takes advantage of unigram and digram probabilities to decide on the characters. That is, in deciding a character, cognizance is taken of:

- * the probability that a proposed character should appear,
- * the probability that the proposed character should appear, given the character that is observed,
- * the probability that the proposed character should appear, given the character that precedes it (which was decided upon already), and perhaps
- * the probability that the proposed character should appear, given the character that succeeds it.

[0057] Over and above the recognition process, in accordance with the principles of this invention there is still room for leaving a character as an "unrecognized box" if, for example, the overall probability measure derived from the above equations is determined to be below a preselected threshold. In such a case, the OCR process 303 is allowed to translate such characters to a number of "possible" characters. Those "possible" characters have a probability measure associated with them, and information is used in the course of the search process. For example, a word such as "London" might be unclear and the recognizer may come back with the following:

"L"
"o"
"n" 70% , "m" 23%
"d"
"c" 46% "o" 68%, "e" 18%
"n" 50% , "m" 33%

By maintaining the collections of possibilities rather than deciding on a word of moderate overall probability measure, a user who wishes to search for, say, the name "Lomdem" would be offered the opportunity to inform the system that, perhaps, the word in question is "London" or, indeed, "Lomdem".

[0058] One of the advantages of this invention is that in spite of the failure by the OCR block to recognize characters that it perhaps ought to have recognized, displays to the user are not hindered by such failures. The user sees the scanned image and, as far as the user knows or cares, all of the characters have been properly recognized by the OCR block.

5 [0059] Another advantage of this invention is that even if the analysis of the displayed image is poor, it is the complete scanned image that is stored and, consequently, the output (both to the screen and/or the printer) may be in full resolution of the scanned image. Reductions may reduce the resolution to that of the screen or the printer, and enlargements may employ conventional interpolation techniques to provide enhanced resolution (at times it may be only a perception of enhanced resolution but that, too, has advantages).

10 [0060] A major advantage of this invention is that searching is performed in the second and third planes of the invention. This provides an effective search mechanism for what appears to be the original images.

[0061] Although there, unbeknownst to the user, instances of the "unrecognized box" entities may exist that might hinder the search process, in accordance with yet another advantage of our invention, the search algorithm evaluates the probability that a data portion (e.g., a word) may represent a match. When that probability is higher than a preselected threshold, the user is informed that a possible match has occurred, the data segment with a possible match is displayed, the unrecognizable portion is highlighted and the user is requested to confirm the match determination. If the user confirms, then the "unrecognized box" entities are replaced with the appropriate characters. If the user determines that the highlighted data portion (e.g., word) does not correspond to a match, the user is given an opportunity to specify what that data portion should be.

20 [0062] The procedure described above basically allows the data generation process to proceed without detours to resolve ambiguities in the OCR process. The user is consulted only when resolution of an ambiguity is in his immediate interest. This is still another advantage of this invention, since it permits a relatively quick means for populating the database with information.

25 [0063] Another advantage of this invention is that improvements in recognition processes (both contextual and otherwise) can be applied to the data within the FIG. 1 system even at a later time. This capability stems directly from the fact that the original scanned data is not discarded. Hence, it can be reprocessed.

[0064] One major goal of this invention clearly is to create a framework for an electronic library. To this end, in the course of discovering this invention and developing its capabilities, numerous technical journals have been scanned into an experimental system. A user of such a system is given the option to either specify a particular journal or to see in a single image all of the journals that are contained in the database. When choosing the latter option, a reduced replica of the scanned images of the first page of the latest issue of the journals are arranged in an array and displayed to the user. The first page of these journals typically contains the journal's logo, date, and other information, such as an indication of the primary topic of that issue. By pointing to the reduced image of a particular journal, the user selects the journal and the particular issue. At that point, a nonreduced image of the journal's first page appears and the user can then request other information relative to that issue, such as the table of contents of the issue (if that is not already contained on the first page) or particular articles. Searching through other issues of that technical journal can also be carried out.

[0065] To demonstrate the versatility of this invention, some patents have also been scanned into the experimental system. The patents have been scanned in their entireties to form a completely searchable database that is capable of displaying to the user images of all pages. The system is also capable of also displaying just the first page of each patent; which typically includes text and a drawing. This is a very powerful capability, because those who wish to carry out novelty searches typically wish to identify patents where certain key words are present in the text. Having identified a number of patents, they typically wish to view the first page of those patents, as an efficient way to reduce the number of patents that need to be studied carefully.

40 [0066] One important use for this invention may be found in connection with litigation. The discovery process in litigation often results in substantial numbers of documents that are delivered by one party to the other. Those documents need to be analyzed and indexed, if they are to serve the needs of the receiving party. Scanning those documents and storing the images with an aid of an OCR system would be very useful. However, those documents often contain handwritten information which existing OCR systems cannot handle. By employing the principles of this invention, the true document images may be stored and, to the extent possible, the OCR recognizable entities are stored in the translated images, and can thereafter be searched. Another problem that discovery documents present relates to the identification of information categories. Whereas, with "form letters" and the like, domain-dependent information can be obtained or derived, with many other documents that is not the case. One solution, in accordance with the principles of this invention, is to create pseudo-domain-dependent information by marking color-highlighted blocks with commercially available wide felt-tip pens having different transparent colors. The area surrounding a date can be marked, for example, with a red highlight; the area surrounding the author's name can be marked, for example, with a yellow highlight; etc. This requires the scanning equipment to be sensitive to colors, of course; but this capability is available in commercial scanners.

Claims

1. A computer implemented method for accessing at least one information segment comprising the steps of:
 - 5 selecting in response to a user request at least one stored image information segment including at least a portion of information associated with textual content; and
displaying at least a portion of the at least one selected stored image information segment along with image content directed to the portion of information associated with the textual content; **CHARACTERISED IN THAT**
10 the stored image information segment comprises a first plane containing a digitized representation of the scanned segment, a second plane containing elemental information comprising textual content that is found in the digitized image and a third plane containing macro information which identifies groupings of elemental information in the first and second planes.
 - 15 2. The method of claim 1, further comprising the steps of:
 - 15 storing at least one scanned image information segment including the at least one portion associated with the textual content to produce the at least one stored image information segment; and
processing the at least one stored image information segment to derive the image content corresponding to the textual content in the at least one scanned image information segment.
 - 20 3. A method according to either one of claims 1 or 2, further comprising displaying correlation information that correlates the stored image information segment to the corresponding image content.
 - 25 4. A method according to any one of claims 1 to 3, further comprising displaying at least one link indicator representing a link between the image content and at least one other stored image information segment.
 - 30 5. A method according to claim 4, further comprising selecting the displayed at least one link indicator and displaying the at least one other stored image information segment.
 - 35 6. A method according to any one of claims 1 to 5, further comprising a step of editing said image content.
 7. A method according to any one of claims 1 to 6, wherein the step of selecting is based on the image content of the at least one stored image information segment.
 - 35 8. A computer implemented method for storing information segments comprising the steps of:
 - 35 storing at least one scanned image information segment including at least a portion of textual content;
processing the at least one scanned image information segment to derive image content associated with the portion of textual content; and
40 storing the image content together with correlation information that correlates the image content with the at least one scanned image information segment;

CHARACTERISED IN THAT

the stored information segment comprises a first plane containing a digitized representation of the scanned segment, a second plane containing elemental information comprising textual content that is found in the digitized image and a third plane containing macro information which identifies groupings of elemental information in the first and second planes.

- 45 9. A method according to claim 8, wherein the processing step employs an optical character recognition step to derive the image content.
- 50 10. A method according to either one of claims 8 or 9, wherein the processing step detects a significance of a particular portion of the image content.

Patentansprüche

- 55 1. Durch Computer implementiertes Verfahren zum Zugreifen auf mindestens ein Informationssegment, mit den fol-

genden Schritten:

als Reaktion auf eine Benutzeranforderung, Auswählen mindestens eines gespeicherten Bildinformationssegments, das mindestens einen Teil von Informationen enthält, die Textinhalt zugeordnet sind; und

5 Anzeigen mindestens eines Teils des mindestens einen gespeicherten Bildinformationssegments zusammen mit Bildinhalt bezüglich des dem Textinhalt zugeordneten Informationsteils; **dadurch gekennzeichnet, daß**

10 das gespeicherte Bildinformationssegment folgendes umfaßt: eine erste Ebene, die eine digitalisierte Darstellung des gescannten Segments enthält, eine zweite Ebene, die elementare, Textinhalt umfassende Informationen enthält, die in dem digitalisierten Bild gefunden werden, und eine dritte Ebene, die Makroinformationen enthält, die Gruppierungen von elementaren Informationen in der ersten und der zweiten Ebene identifizieren.

2. Verfahren nach Anspruch 1, weiterhin mit den folgenden Schritten:

15 Speichern mindestens eines gescannten Bildinformationssegments, einschließlich des mindestens einen Teils, der dem Textinhalt zugeordnet ist, um das mindestens eine gespeicherte Bildinformationssegment zu erzeugen; und

20 Verarbeiten des mindestens einen gespeicherten Bildinformationssegments, um den Bildinhalt abzuleiten, der dem Textinhalt in dem mindestens einen gescannten Bildinformationssegment entspricht.

3. Verfahren nach einem der Ansprüche 1 oder 2, bei dem weiterhin Korrelationsinformationen angezeigt werden, die das gespeicherte Bildinformationssegment mit dem entsprechenden Bildinhalt korrelieren.

25 4. Verfahren nach einem der Ansprüche 1 bis 3, bei dem weiterhin mindestens eine Verknüpfungsanzeige angezeigt wird, die eine Verknüpfung zwischen dem Bildinhalt und mindestens einem weiteren gespeicherten Bildinformationssegment darstellt.

30 5. Verfahren nach Anspruch 4, bei dem weiterhin die angezeigte mindestens eine Verknüpfungsanzeige ausgewählt und das mindestens eine weitere gespeicherte Bildinformationssegment angezeigt wird.

6. Verfahren nach einem der Ansprüche 1 bis 5, weiterhin mit dem Schritt des Editierens des Bildinhalts.

35 7. Verfahren nach einem der Ansprüche 1 bis 6, wobei der Schritt des Auswählens auf dem Bildinhalt des mindestens einen gespeicherten Bildinformationssegments basiert.

8. Durch Computer implementiertes Verfahren zum Speichern von Informationssegmenten, mit den folgenden Schritten:

40 Speichern mindestens eines gescannten Bildinformationssegments, das mindestens einen Teil mit Textinhalt enthält;

45 Verarbeiten des mindestens einen gescannten Bildinformationssegments, um dem Textinhaltteil zugeordneten Bildinhalt abzuleiten; und

Speichern des Bildinhalts zusammen mit Korrelationsinformationen, die den Bildinhalt mit dem mindestens einen gescannten Bildinformationssegment korrelieren;

50 **dadurch gekennzeichnet, daß**

das gespeicherte Informationssegment folgendes umfaßt: eine erste Ebene, die eine digitalisierte Darstellung des gescannten Segments enthält, eine zweite Ebene, die elementare, Textinhalt umfassende Informationen enthält, die in dem digitalisierten Bild gefunden werden, und eine dritte Ebene, die Makroinformationen enthält, die Gruppierungen von elementaren Informationen in der ersten und der zweiten Ebene identifizieren.

55 9. Verfahren nach Anspruch 8, wobei der Verarbeitungsschritt einen Schritt der optischen Zeichenerkennung umfaßt, um den Bildinhalt abzuleiten.

10. Verfahren nach einem der Ansprüche 8 oder 9, wobei der Verarbeitungsschritt eine Signifikanz eines bestimmten Teils des Bildinhalts erkennt.

5 **Revendications**

1. Procédé mis en oeuvre par ordinateur pour accéder à au moins un segment d'information comprenant les étapes de :

10 sélection en réponse à une demande d'utilisateur d'au moins un segment d'information d'image mémorisé comportant au moins une partie d'information associée à un contenu textuel ; et affichage d'au moins une partie du au moins un segment d'information d'image mémorisé sélectionné avec un contenu d'image se rapportant à la partie d'information associée au contenu textuel ; **CARACTERISE EN CE QUE**

15 le segment d'information d'image mémorisé comprend un premier plan contenant une représentation numérisée du segment balayé, un deuxième plan contenant des informations élémentaires comprenant un contenu textuel trouvées dans l'image numérisée et un troisième plan contenant des macro-informations qui identifient des groupements d'informations élémentaires dans les premier et deuxième plans.

20 2. Procédé selon la revendication 1, comprenant en outre les étapes de :

25 mémorisation d'au moins un segment d'information d'image balayé comportant la au moins une partie associée au contenu textuel afin de produire le au moins un segment d'information d'image mémorisé ; et traitement du au moins un segment d'information d'image mémorisé afin de dériver le contenu d'image correspondant au contenu textuel dans le au moins un segment d'information d'image.

30 3. Procédé selon l'une ou l'autre des revendications 1 ou 2, comprenant en outre l'affichage d'informations de corrélation qui mettent en corrélation le segment d'information d'image mémorisé et le contenu d'image correspondant.

35 4. Procédé selon l'une quelconque des revendications 1 à 3, comprenant en outre l'affichage d'au moins un indicateur de lien représentant un lien entre le contenu d'image et au moins un autre segment d'information d'image mémorisé.

5. Procédé selon la revendication 4, comprenant en outre la sélection du au moins un indicateur de lien affiché et l'affichage du au moins un autre segment d'information d'image mémorisé.

6. Procédé selon l'une quelconque des revendications 1 à 5, comprenant en outre une étape d'édition dudit contenu d'image.

40 7. Procédé selon l'une quelconque des revendications 1 à 6, dans lequel l'étape de sélection est basée sur le contenu d'image du au moins un segment d'information d'image mémorisé.

8. Procédé mis en oeuvre par ordinateur pour mémoriser des segments d'information comprenant les étapes de :

45 mémorisation d'au moins un segment d'information d'image balayé comportant au moins une partie de contenu textuel ; traitement du au moins un segment d'information d'image mémorisé afin de dériver un contenu d'image associé à la partie de contenu textuel ; et mémorisation du contenu d'image avec des informations de corrélation qui mettent en corrélation le contenu d'image avec le au moins un segment d'information d'image balayé ;

CARACTERISE EN CE QUE

55 le segment d'information mémorisé comprend un premier plan contenant une représentation numérisée du segment balayé, un deuxième plan contenant des informations élémentaires comprenant un contenu textuel trouvées dans l'image numérisée et un troisième plan contenant des macro-informations qui identifient des groupements d'informations élémentaires dans les premier et deuxième plans.

9. Procédé selon la revendication 8, dans lequel l'étape de traitement emploie une étape de reconnaissance optique

EP 0 539 106 B1

de caractères afin de dériver le contenu d'image.

10. Procédé selon l'une ou l'autre des revendications 8 ou 9, dans lequel l'étape de traitement détecte une importance d'une partie particulière du contenu d'image.

5

10

15

20

25

30

35

40

45

50

55

FIG. 1

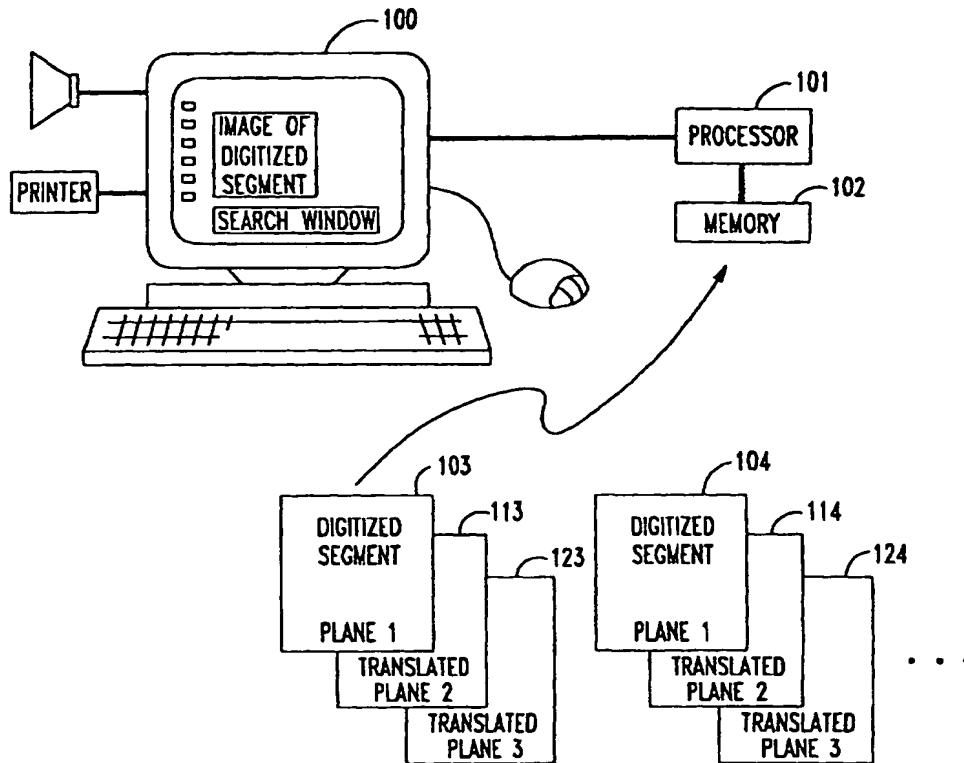


FIG. 2

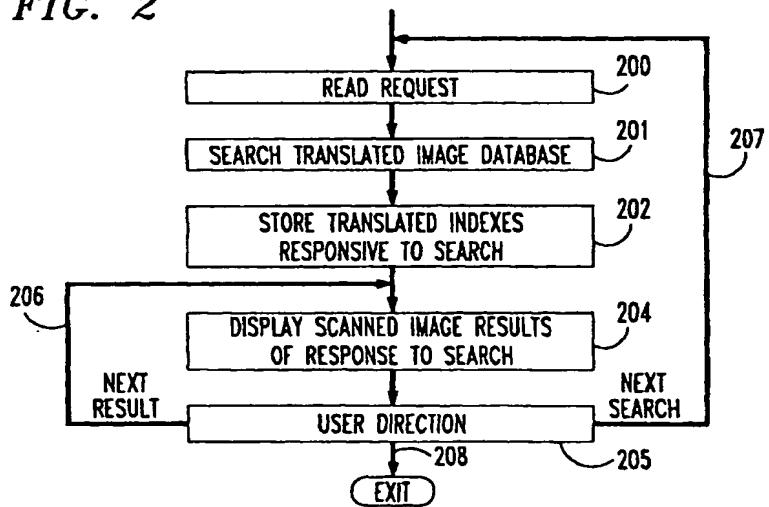


FIG. 3

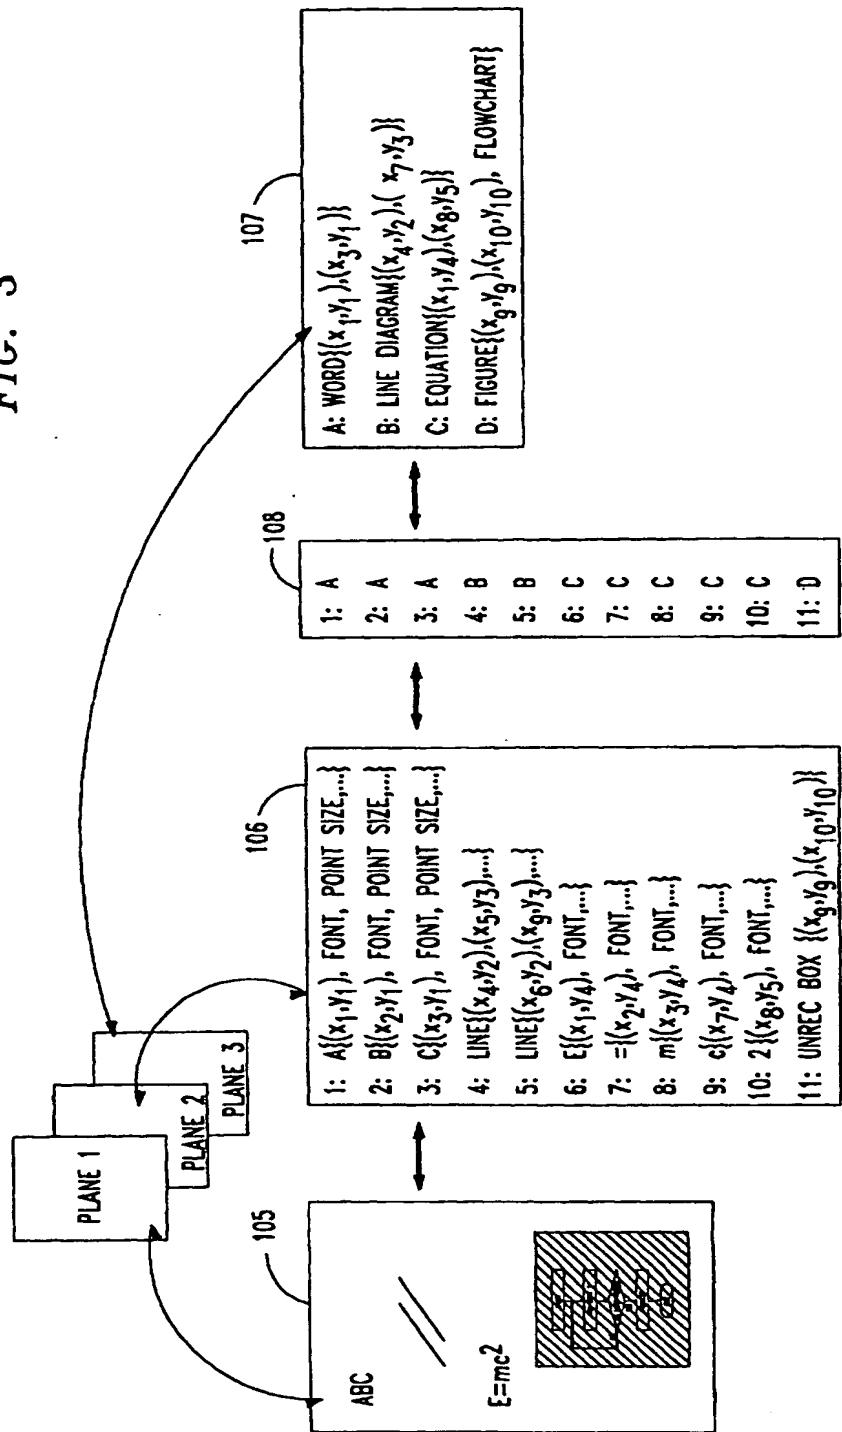


FIG. 4

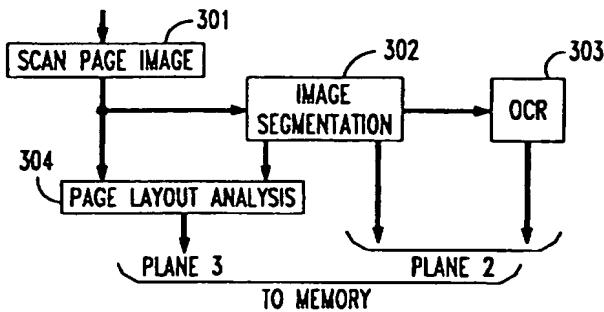


FIG. 5

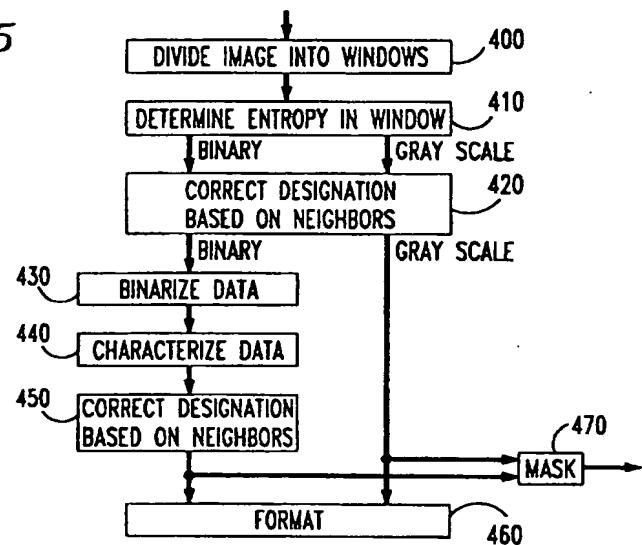


FIG. 6

